IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

HP's sx1000 Chipset: Innovation Atop Standardization

Sponsored by: Hewlett-Packard

Vernon Turner Shane Rau

February 2003

EXECUTIVE SUMMARY

Hewlett-Packard's new super scalable processor chipset, sx1000, is designed to meet emerging feature sets and required scalability in the server marketplace and to differentiate HP servers that use the Itanium processor from other servers using Itanium. The chipset, which consists of five types of application-specific integrated circuits (ASICs), supports dual-processor modules (i.e., two processors installed in a single socket) and provides higher CPU, memory, and I/O bandwidth and greater memory capacity. The chipset supports both the Itanium 2 processor (starting with the Madison version) and the PA-RISC 8800 processor, thus providing a hardware migration path from PA-RISC to Itanium processors.

The sx1000 chipset provides essential technology for HP's server building blocks, which are called cells. Each cell has four sockets that can be populated with single- or dual-processor modules. Servers assembled from these cells will scale to 128 Itanium 2 or PA-RISC 8800 processors. Hard partitions formed by these cells will allow the resources of a single HP server to be divided among many enterprise workloads and to provide different operating environments (e.g., HP-UX, Linux, Microsoft Windows Server 2003, OpenVMS) simultaneously.

IDC has identified five trends in the server marketplace:

- System suppliers are competing to provide scalable server architectures that use standard components, particularly the Intel Itanium processor.
- When consumers consolidate server resources to reduce cost, they are shopping for more powerful servers that can be partitioned.
- The ability to provision IT software and services, including server resources, is a step in the evolution of utility computing, the next-generation datacenter architecture.
- IT consumers expect servers to support Unix, Linux, and Microsoft operating environments.
- Consumers and suppliers believe that standardization of some components will lead to lower costs and to value being added atop those standard components.

HP's sx1000 chipset functionality addresses IDC's server marketplace trends, albeit at a deep technical level. By providing memory and processor buses with higher bandwidth, the sx1000 chipset ensures that more powerful processors are provided data at the high rates they need. Fast, low-level error correction further enhances the server's overall availability. Support for new dual-processor modules extends scalability for HP servers from 64-way to 128-way products.

Readers who wonder whether standardization on the Intel Itanium processor will lead to a commodity market of servers with equivalent performance should look carefully at what HP is doing. Innovation continues to fuel competition in the server marketplace as HP and other vendors shift their investments from processor technology to the many other technologies necessary to produce a compute server. HP intends the sx1000 chipset to be one of many differentiators resting atop an emerging standard component.

INTRODUCTION

IDC believes that five key trends describe the most important activities in the server marketplace. These trends are a mixed collection of technology advances, supplier initiatives, and shifting requirements from IT consumers, as discussed below:

- **Vendors are extending their product lines by providing standardized, but highly scalable, systems.** System suppliers provide a continuum of offerings from single-processor servers to multiprocessor servers with 2 to 64 or more processors working in parallel. Clusters of servers offer even greater capacity along with redundancy for improved availability.

- **IT planners are consolidating resources to reduce cost throughout IT operations.** Server consolidation, aimed at reducing the high maintenance costs associated with multitudes of smaller servers, requires more powerful machines that can be partitioned to serve workgroups and workloads of different sizes.

- **IT futurists are looking ahead to the next model for IT systems — the utility computing model.** A utility model assumes that consumers of IT are not interested in owning servers and storage devices but rather in obtaining processing and storage services on an as-needed basis. System consolidation is a step in the direction of utility computing.

- **IT managers have come to believe that a small collection of operating environments will coexist in the enterprise.** In addition to enclaves of proprietary operating systems provided by system vendors, IDC believes that any server claiming to conform to industry standards will need to support three major operating environments: Unix, Linux, and Microsoft Windows Server 2003. Further, if servers can do so concurrently in different partitions, so much the better.

- **Standardization continues to play a critical role in the thinking of both suppliers and consumers of server technology.** When system suppliers use industry-standard components, there are greater economies of scale for component producers, leading to lower prices. Moreover, system suppliers that use standard components can focus more of their resources on system-level features that meet customer requirements. In the enterprise server and high-performance computing markets, standardization is all about moving to the Itanium processor architecture and ecosystem.

HP's server strategy is congruent with IDC's view of server market trends. HP's Business Critical Systems group is extending its scalable line of products to support server consolidation and, eventually, a utility model for delivering IT services. HP's servers host the three major operating environments that customers value, and HP has for many years been planning for a transition from its own processors to the Itanium processor family.

This white paper focuses on the sx1000 chipset that HP has announced for use in its mid- and upper-range servers. IDC views the sx1000 chipset as an example of how technology investment in conjunction with the use of standard components results in distinguishing features and enhanced capabilities for IT customers.

THE HP PROCESSOR CHIPSET SX1000

HP has released the super scalable processor chipset sx1000, a new chipset for its server products. Like the prior chipset, code-named Yosemite, the sx1000 chipset consists of five types of application-specific integrated circuits (ASICs) that surround the server's processors and provide connectivity to system resources.

BENEFITS OF THE CELL DESIGN

HP servers using the sx1000 chipset are based on building blocks called cells. Cells enable the deployment of hard partitions (which HP calls nPars). Hard partitions, in turn, enable the isolation of failed cells so that their workload can be offloaded to other cells and the failed cells can be replaced easily. Partitions also enable multiple instances of operating systems to be deployed on a single server for server consolidation. Operating environments can vary in each hard partition, enabling a single server to run HP-UX, Linux, Microsoft Server 2003, and OpenVMS, all at the same time. Partitioning allows computer resources to be dynamically reallocated, which is the basis of HP's vision for utility computing.

CELLS AND THE SX1000 CHIPSET

Each cell consists primarily of a printed circuit board supporting four processor sockets and nine ASICs — a cell controller and eight memory buffer chips. The cell controller and memory buffer ASICs are two of the five chip types that compose the sx1000 chipset. Into each socket goes a processor module that, as noted later, may contain one or two processors. The smallest system using the sx1000 chipset will support up to two cells (eight sockets). The largest system will support up to 16 cells (64 sockets).

As Figure 1 shows, each cell controller, which is the key sx1000 chip, routes data and signaling to the cell's four CPU modules and provides control and buffering for four independent memory subsystems. The cell controller connects to other cells through a third sx1000 ASIC, the crossbar switch. The cell controller also connects to an I/O backplane board, on which reside the remaining ASICs, a PCI-X system bus adapter, and PCI-X host bridges that interface to PCI and PCI-X cards. The PCI-X system bus adapter may be connected to the central controller by printed circuit or by cable.

Cells and ASIC crossbar switches are assembled into 8-socket, 16-socket, and 64-socket systems, as shown in Figure 2. Note that a crossbar ASIC is not required when joining two cells into an 8-socket system.

FIGURE 1

HP SX1000 CELL AND I/O BACKPLANE BLOCK DIAGRAM

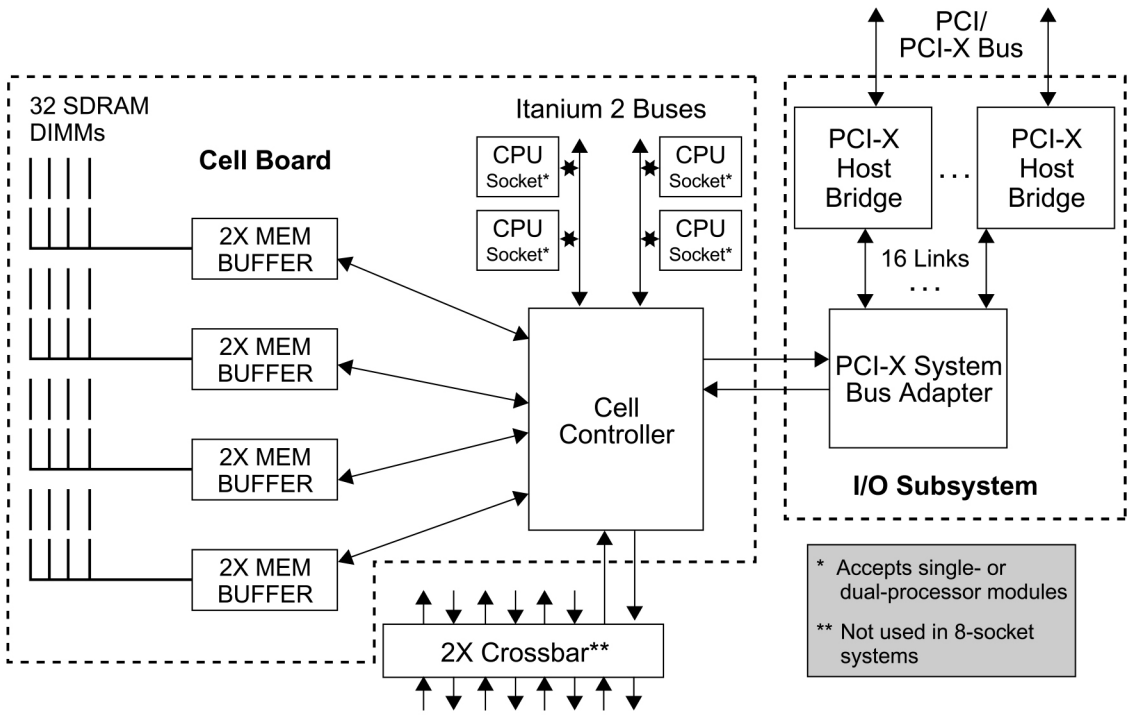

Source: Hewlett-Packard, 2003

DESIGN OBJECTIVES FOR THE HP SX1000 CHIPSET

HP engineers designed the sx1000 chipset to meet three key objectives:

- **The sx1000 chipset will support both Itanium processors and PA-RISC processors.** Specifically, the sx1000 chipset will support the Itanium 2 processor (code-named Madison) or the PA-RISC 8800 processor. However, cells with Itanium 2 processors and PA-RISC processors will not be mixed in the same system. An important benefit of supporting Itanium processors is that several operating systems have been or will be ported to Itanium: HP-UX, Linux, Microsoft Windows Server 2003, and OpenVMS.

- **The sx1000 chipset will support both single- and dual-processor CPU modules.** A dual-processor module provides two CPUs in a package that will plug into a single socket on the CPU bus. The PA-8800, available in late 2003, will be the first dual-core processor in the PA-RISC processor line. HP will also introduce the hp mx2 dual-processor module in early 2004. The hp mx2 contains a daughter card with two Madison processors and a 32MB L4 cache. High processor bus and memory bandwidth will support the elevated data rates required by these powerful processors.

- **The sx1000 chipset will support cell and I/O subsystem upgrades that use the server's current memory, system backplane, and PCI cards.** However, all computing cells in a server will need to be swapped when transitioning from PA-RISC 8700 and older processors to Itanium 2 or PA-RISC 8800 processors. Also, high-end customers (i.e., SuperDome customers) may choose to keep their old PCI-only I/O subsystem or upgrade to a PCI-X subsystem. Customer investment in backplane hardware, PCI cards, memory, and time spent connecting the server to datacenter infrastructure is preserved.

FIGURE 2

HP CELL ARCHITECTURE OVERVIEW

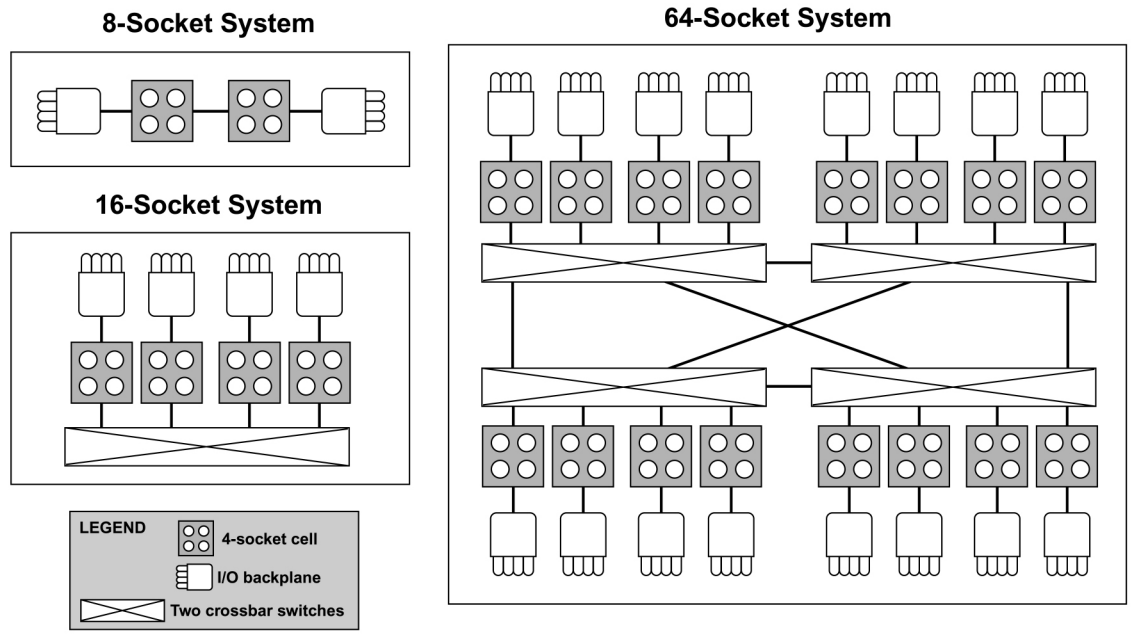

Source: Hewlett-Packard, 2003

SX1000 CHIPSET FUNCTIONALITY

Four of the five types of sx1000 chips offer significant new functionality. Although the crossbar chip remains unchanged from the previous chipset, the cell controller, memory buffer, PCI-X system bus adapter, and PCI-X host bridge offer higher performance and new functionality when compared with the Yosemite chipset that powers today's HP cell-based servers.

At the block level, Figure 3 shows the contrast between the Yosemite chipset and the new sx1000 chipset. In addition to supporting both Itanium and PA-RISC processors, the new chipset provides twice the bandwidth when CPU sockets access memory on their local cell. The chipset's interface to the CPUs also runs at 1.6-times greater bandwidth than the current chipset. The sx1000 chipset's PCI-X system bus adapter (SBA) and host bridge are compatible with both the PCI interface and the new, faster PCI-X interface.

FIGURE 3

HP YOSEMITE AND SX1000 CHIPSET BLOCK DIAGRAMS

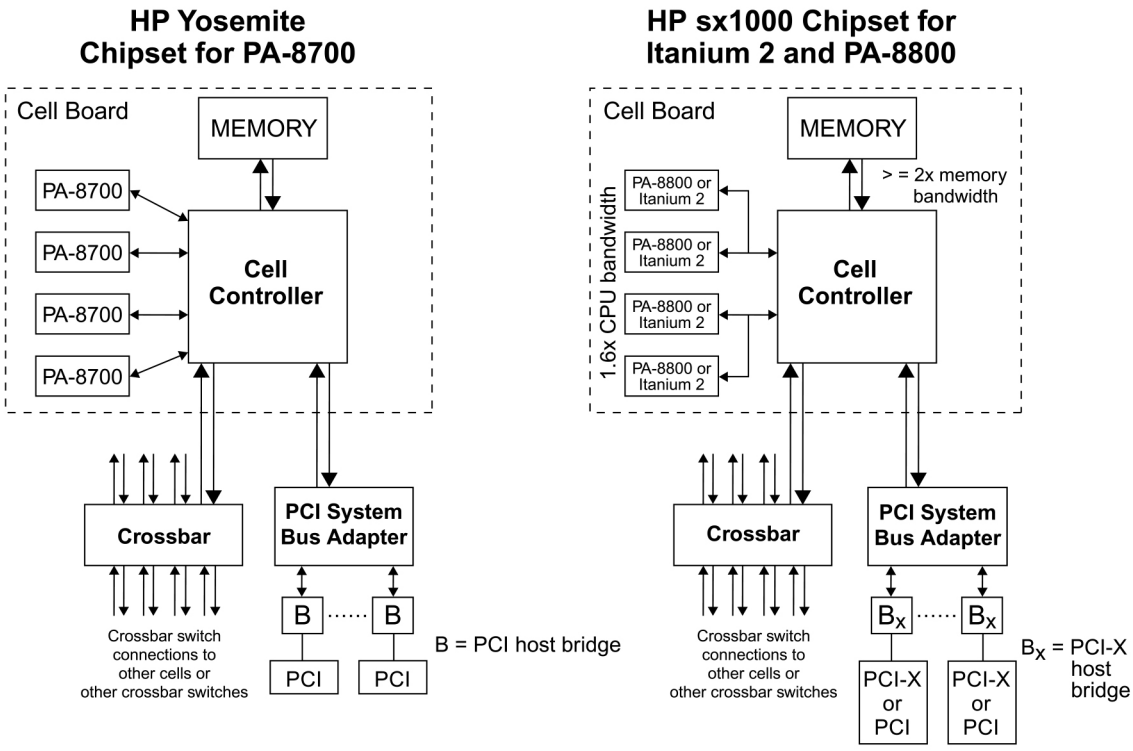

Source: Hewlett-Packard, 2003

CELL CONTROLLER

To support next-generation Itanium 2 (code-named Madison) and PA-RISC dual-processor modules, the sx1000 chipset's cell controller provides two processor buses with a combined peak bandwidth of 12.8GBps, which is 1.6-times greater than the 8GBps combined bandwidth of the previous chipset's processor buses.

The cell controller manages four independent memory subsystems with a combined peak bandwidth of 16GBps for each cell, which is twice the bandwidth of the previous chipset's memory subsystem. The new sx1000 cell controller is backward compatible with the previous PCI I/O subsystem as well as with the newer PCI-X I/O subsystem. A glossary of other cell controller features is given in Table 1.

MEMORY BUFFER

A pair of memory buffer ASICs in each of the four memory subsystems has a 288-bit bus connection to eight SDRAM DIMMs, thus supporting a total of 32 DIMM slots per cell. With today's 2GB DIMMS, 64GB of memory can be placed in each cell. With tomorrow's 4GB DIMMs, the maximum memory capacity will be 128GB per cell and 2TB per server.

TABLE 1

GLOSSARY OF HP CHIPSET SX1000 FEATURES

Chip	Feature	Function
Cell controller	ECC	Special error-correcting code (ECC) circuitry that checks the integrity of data as it passes in and out of memory or on the sx1000 chipset's cell controller interfaces: CPU, crossbar, and I/O
	Directory-based cache coherence	Ensures that data in processor caches is consistent with data in each cell's four local memory subsystems (Directory information is stored in the main memory, along with data, to record the CPUs that are caching data.)
	Chip-kill	Enhanced ECC circuitry that corrects memory errors when an entire memory chip fails
	Interleaved or localized memory support	Maximizes system flexibility by allowing the operating system to treat the memory subsystems across cells uniformly or independently (Interleaving memory across cells allows for scalability and performance. Localizing memory allows for data specific to one cell's processor to be stored on that cell. Windows Server 2003 and HP-UX can take advantage of this feature to improve performance.)
Memory buffer	Built-in directory update circuitry	Mechanism to quickly update information in the cache coherence directory stored in memory (This minimizes memory bandwidth overhead associated with directory-based coherence.)
PCI-X system bus adapter	Caches	Two 8KB caches speed DMA transfers by PCI cards
	Translation buffer	Provides 32-bit PCI cards access to all memory (up to 2TB)
PCI-X host bridge	Support for IO SAPIC interrupt protocol	Mechanism for directing PCI interrupts to CPUs in the system

Source: Hewlett-Packard, 2003

PCI-X SYSTEM BUS ADAPTER

The sx1000 chipset's PCI-X system bus adapter provides an upgrade from the previous chipset's PCI system bus adapter. It offers 16 10-bit high-speed links, and each 10-bit link running at 266MHz transfers at dual data rate. A pair of these high-speed links can be connected to a PCI-X host bridge to support full-speed PCI-X data rates (about 1GBps). Only one link can be connected to a host bridge, if PCI data rates are sufficient for the application (about 500MBps). There are two 8KB caches in the system bus adapter that participate in cache coherency. Data can be fetched into these caches for faster access by PCI cards.

An address translation mechanism in the sx1000 chipset's I/O system bus adapter allows 32-bit PCI cards to access any memory in the whole system directly (i.e., up to 2TB). Although the speed of the links to the PCI-X host bridges has been doubled with the sx1000 chipset, the other features mentioned existed in the previous chipset.

PCI-X HOST BRIDGE

The PCI-X host bridge provides an upgrade from the previous chipset's PCI host bridge. 3.3V PCI and PCI-X cards are supported, including the ability to add and replace cards while the system is operating (i.e., online addition and replacement [OLA/R]). The PCI-X host bridge provides two 10-bit high-speed links, each running at 532 megatransfers per second, to the PCI-X system bus adapter to support increasing bandwidth requirements for high-speed network connections. Only one link may be connected if PCI bandwidth, not full PCI-X bandwidth, is required. The PCI-X host bridge also contains an IO SAPIC interrupt controller.

HP PROCESSOR CHIPSET SX1000 TECHNOLOGY

The sx1000 chipset's technology is quite advanced, as compared with other chipsets in the industry. All five chips are manufactured using 0.18-micron CMOS technology. As shown in Table 2, the cell controller contains 41 million transistors and hosts 1,466 signal pins in a 2,500-pin package.

TABLE 2

HP SX1000 CHIPSET TECHNOLOGY

Chip	Transistors (M)	Signal Pins
Cell controller	41	1,466
Memory buffer	1	248
System bus adapter	14	612
PCI-X host bridge	12	170
Crossbar	21	748

Source: Hewlett-Packard, 2003

SUMMARY OF HP SX1000 CHIPSET FEATURES

Performance improvements for the sx1000 chipset include increases in most key parameters, as shown in Table 3. Key improvements are in processor bus and memory bandwidth, address range, processors per cell, and memory capacity.

TABLE 3

HP SX1000 CHIPSET VERSUS HP YOSEMITE CHIPSET

Feature	Yosemite	sx1000
Processors	PA-8600, PA-8700	PA-8800, Itanium 2
Core frequency	250MHz	250MHz
Processor bus	4 Runway buses, 8 bytes, 250MTps	2 Itanium 2 buses, 16 bytes, 400MTps
Physical address	40 bits	44 bits
Cache line size	64 bytes	128 bytes
CPUs per cell	4	4–8
Peak memory bandwidth per cell	8GBps	16GBps
Memory capacity	128GB per cell, 512GB per system	128GB per cell, 2TB per system
I/O card	PCI	PCI and PCI-X
Operating system support	HP-UX 11i	HP-UX 11i, HP-UX 11i v2, Linux, Microsoft Windows Server 2003
High availability	Hard partitions; online addition/deletion of CPUs, memory, PCI cards; ECC on memory, I/O, and crossbar; chip-kill memory	Hard partitions; online addition/deletion of CPUs, memory, PCI cards; ECC on memory, I/O, crossbar, and processor bus; chip-kill memory

Source: Hewlett-Packard, 2003

HP CHIPSET SX1000 SUMMARY

HP chipset sx1000 provides a uniform cell architecture that supports single- and dual-processor Itanium 2 and PA-RISC modules. This capability will ease the transition from PA-RISC to Itanium 2 at the hardware level. Server upgrades can be accomplished without replacement of the system backplane and I/O backplane (on high-end systems), PCI cards, and memory DIMMs. Improved scalability is provided through support for dual-processor modules, high CPU and memory bandwidth, and a large I/O and memory capacity.

IDC ANALYSIS: OPPORTUNITIES AND CHALLENGES

IDC believes that HP's decision to move to the industry-standard Itanium processor family and HP's early commitment to being a strong partner with Intel in developing Itanium technologies were good decisions. By creating the sx1000 chipset, HP is exploiting its deep understanding of both PA-RISC and Itanium processor architectures. HP chipset sx1000 stands as an early proof point that standardized components can and will be enhanced by auxiliary technologies, and the sx1000 chipset underscores the notion that a computer server's capability depends on sophisticated hardware technologies that are tightly integrated with the processor architecture.

The transition to Itanium processors will raise new challenges for HP, however. Should the adoption rate of Itanium processors be lower than industry expectations, HP will be pressed by its customers and prospects to provide compelling evidence that its Itanium processor–based servers do, in fact, provide greater performance than competitive servers from companies that eschew the Itanium processor. Although the sx1000 chipset lays the groundwork for higher performance on standard server benchmarks, HP needs to develop empirical support indicating that the transition to Itanium processor–based servers results in greater capabilities at the server level of analysis.

Standardization on a costly component such as the server's processor lowers the barriers to entry and broadens the competition for HP, which is a new challenge. Suppliers unable to afford proprietary processor development are now ostensibly on a more level playing field. Value-added functionality is now the focus for HP and other system suppliers previously locked out of the high-performance server market.

HP's installed base, encompassing both legacy HP and legacy Compaq customers, is a valuable asset for HP. Value-added functionality is a familiar competitive weapon for HP. Combining the strengths of HP-UX and Tru64 UNIX and embracing Linux are examples of how HP intends to capitalize on its knowledge of the needs of segments of its customer base.

Ironically, for HP and for the server industry as a whole, servers based on legacy IA-32 processors remain a significant threat to servers based on next-generation Itanium processors. Clusters of 32-bit SMP processors offer a price/performance advantage that is likely to inhibit the transition to 64-bit processors. Although IDC believes that the need for high-performance 64-bit computing will emerge, we forecast that 64-bit systems will be prevalent late in this decade.

HP can meet this IA-32 challenge, however, by providing Itanium processor–based compute servers with improved price/performance ratios and greater sheer performance than servers based on contemporary RISC and IA-32. The sx1000 chipset is certain to be a key component for this new generation of server products.

CONCLUSION

HP intends the new sx1000 chipset to be more than an ordinary chipset upgrade. HP's design principles for the sx1000 chipset were tied explicitly to the company's server strategy, which is to leverage emerging industry-standard Itanium processors with innovative technologies that amplify the processor's performance and its ability to be fielded effectively. Although HP will most certainly face challenges from its competitors, IDC believes that the announcement of the sx1000 chipset is an important competitive event in a new and different world of server design.